arXiv:1603.07364v2 [math.CO] 17 Jan 2017

SPECIAL DIVISORS ON MARKED CHAINS OF CYCLES

NATHAN PFLUEGER

Abstract. We completely describe all Brill-Noether loci on metric graphs consisting of a chain of g cycles with arbitrary edge lengths, generalizing work of Cools, Draisma, Payne, and Robeva. The structure of these loci is determined by displacement tableaux on rectangular partitions, which we define. More generally, we fix a marked point on the rightmost cycle, and completely analyze the loci of divisor classes with specified ramification at the marked point, classifying them using displacement tableaux. Our results give a tropical proof of the generalized Brill-Noether theorem for general marked curves, and serve as a foundation for the analysis of general algebraic curves of fixed gonality.

1. Introduction

The tropical proof of the Brill-Noether theorem [CDPR12] gave a parameterization of the sets $W^r_d(\Gamma)$ of special divisors on a metric graph Γ composed of a chain of cycles with generic edge lengths. The purpose of this paper to to generalize this analysis, in two ways. We use these generalizations to give a tropical proof of the generalized Brill-Noether theorem (for curves with a marked point), and also as a foundation for our paper [Pfl16], which gives applications to the geometry of general curves of fixed gonality.

The first generalization is that we consider *arbitrary* chains of cycles, with any edge lengths. We show that the behavior of the loci of special divisor classes on Γ depends only on a sequence of numbers $\underline{m} = (m_2, m_3, \cdots, m_g)$, easily computed from the edge lengths, called the *torsion profile* of Γ (Definition 1.9). The torsion profile determines whether the chain Γ is Brill-Noether general in the sense of [CDPR12]. The metric graph Γ is called *Brill-Noether general* if $\dim W^r_d(\Gamma) = g - (r+1)(g-d+r)$ whenever this number is nonnegative, and $W^r_d(\Gamma)$ is empty otherwise. The original genericity condition of [CDPR12] is equivalent to saying that $m_i = 0$ or $m_i > 2g-2$.

Theorem 1.1. *A chain of cycles Γ is Brill-Noether general if and only if for each $i \in \{2, 3, \cdots, g-1\}$, either $m_i = 0$ or*

$$m_i > \min(i, g+1-i).$$

The second direction in which we generalize [CDPR12] is by considering special divisors with prescribed ramification at a marked point $w \in \Gamma$, which

Date: January 19, 2017.

we always take to be on the rightmost cycle of the chain. In algebraic geometry, imposing ramification conditions amounts to considering only those linear series that correspond to maps to projective space in which the marked point is inflected in a particular way (see Definition 1.6). Although this generalization is interesting in its own right, it is in fact a necessary ingredient in our arguments even in the ordinary case. We will define, for any metric graph Γ with a marked point w, generalized Brill-Noether loci $W^\lambda(\Gamma, w) \subseteq \mathrm{Pic}^0(\Gamma)$, where λ is a partition.

Here by a *partition* we mean a finite, non-increasing sequence of nonnegative integers, where two such sequences are considered the same if one is obtained by adding a sequence of 0s to the end of the other. We will also identify partitions with their Young diagrams (in French notation), which we in turn regard as subsets of $\mathbf{Z}^2_{>0}$, according to the following convention.

Convention 1.2. Any partition $\lambda = (\lambda_0, \lambda_1, \cdots, \lambda_r)$ (where $\lambda_0 \geqslant \lambda_1 \geqslant \cdots \geqslant \lambda_r$) will be identified with the subset

$$\{(x, y) \in \mathbf{Z}^2_{>0} : \ 1 \leqslant x \leqslant \lambda_{y-1}, \ 1 \leqslant y \leqslant r+1\},$$

which we will refer to as the Young diagram of λ. We denote the number of elements in this set (i.e. the sum of the elements λ_i) by $|\lambda|$.

The link between the generalized Brill-Noether loci $W^\lambda(\Gamma, w)$ and the ordinary Brill-Noether loci is that $W^r_d(\Gamma)$ is isomorphic to $W^\lambda(\Gamma, w)$, where λ is a *rectangular* partition of height $r+1$ and width $g-d+r$, and w is any marked point. The isomorphism $W^\lambda(\Gamma, w) \to W^r_d(\Gamma)$ is given by $[D] \mapsto [D + d \cdot w]$.

$\lambda =$ 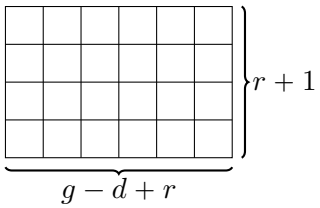

$$W^r_d(\Gamma) \cong W^\lambda(\Gamma, w)$$

The analog of the Brill-Noether number is $g - |\lambda|$. Note in particular that when λ is the rectangle above, this is the usual Brill-Noether number $g - (r+1)(g-d+r)$. We call a *marked* metric graph (Γ, w) *Brill-Noether general* if for all partitions λ, $\dim W^\lambda(\Gamma, w) = g - |\lambda|$ if this number is nonnegative, and $W^\lambda(\Gamma, w)$ is empty otherwise. We obtain the following marked-point version of Theorem 1.1.

Theorem 1.3. *Let (Γ, w) be a marked chain of cycles with w on the rightmost cycle and torsion profile $\underline{m} = (m_2, \cdots, m_g)$. Then (Γ, w) is Brill-Noether general if and only if for each $i \in \{2, 3, \cdots, g\}$, either $m_i = 0$ or $m_i > i$.*

Our main result is an explicit parameterization of the loci $W^\lambda(\Gamma, w)$, when Γ is any chain of cycles of genus g (with arbitrary edge lengths), w is a

marked point on the rightmost cycle, and λ is any partition. This parameterization is based on combinatorial objects called $\underline{m}$-*displacement tableaux* (Definition 2.1), which are Young tableaux on λ with certain constraints depending on the torsion profile $\underline{m}$. We write $t \vdash_{\underline{m}} \lambda$ to indicate that t is an $\underline{m}$-displacement tableau on the partition λ. Each such tableau t defines a locus $\mathbf{T}(t) \subseteq \mathrm{Pic}^0(\Gamma)$, homeomorphic to a torus of dimension equal to g minus the number of symbols appearing in t (Definition 3.5). In particular, $\dim \mathbf{T}(t) \geqslant g - |\lambda|$, with equality if and only if λ has no repeated symbols.

Theorem 1.4. *For any chain Γ of torsion profile $\underline{m}$ and any partition λ,*

$$W^\lambda(\Gamma, w) = \bigcup_{t \vdash_{\underline{m}} \lambda} \boldsymbol{T}(t).$$

A consequence of this (Corollary 3.7) is that we can compute the dimensions of the loci $W^\lambda(\Gamma, w)$ by determining the minimum number of distinct symbols in an $\underline{m}$-displacement tableau on λ. Together with semicontinuity results for tropicalization of algebraic curves (summarized in Section 5), we can deduce upper bounds on the dimensions of Brill-Noether varieties of (marked or unmarked) algebraic curves. Most importantly, choosing *special* edge lengths in a chain of cycles opens new applications of the theory of linear series on metric graphs, as in our [Pfl16]. We also obtain a tropical proof of the generalized Brill-Noether theorem for algebraic curves with one marked point, originally proved by Eisenbud and Harris [EH86, Theorem 4.5].

Theorem 1.5 (Generalized Brill-Noether theorem)**.** *Let (C, p) be a general marked algebraic curve of genus g over an algebraically closed field of any characteristic. Let r, d be positive integers such that $g - d + r \geqslant 0$, and let $\alpha = (\alpha_0, \alpha_1, \cdots, \alpha_r)$ be a nondecreasing sequence of nonnegative integers. Then the variety $W_d^{r,\alpha}(C, p) \subseteq \mathrm{Pic}^d(C)$ of line bundles of rank at least r and ramification at least α at the point p (see Definition 1.6) is nonempty if and only if the adjusted Brill-Noether number*

$$\rho(g, d, r, \alpha) = g - (r+1)(g - d + r) - \sum_{i=0}^{r} \alpha_i$$

is nonnegative. If nonempty, this locus has dimension exactly $\rho(g, d, r, \alpha)$.

This theorem will follow from Corollary 5.3. We now define and give the geometric interpretation of generalized Brill-Noether loci, in both the algebraic and tropical contexts.

1.1. **Brill-Noether loci of marked algebraic curves.** Let C be a smooth projective algebraic curve of genus g, and let d, r be nonnegative integers with $r \geqslant d - g$. Brill-Noether theory concerns the geometry of the schemes $W_d^r(C)$ parameterizing degree d line bundles L on C such that $h^0(L) \geqslant r+1$. The reason that we require $r \geqslant d - g$ is that the Riemann-Roch formula guarantees that $h^0(L) \geqslant (d - g) + 1$, regardless of the line bundle. If a

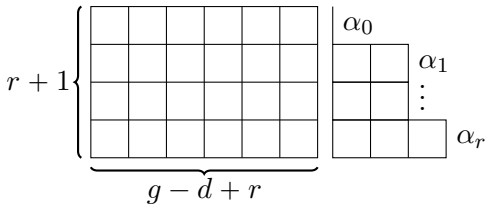

FIGURE 1. The partition λ associated to the data g, d, r, α. In this example, $r = 3$, $d = g - 3$, and $\alpha = (0, 2, 2, 3)$.

marked point $p \in C$ is chosen, one can also impose ramification data at p. The *vanishing orders* of a line bundle L of rank r (that is, $h^0(L) = r + 1$) at the point p are the integers $a_0 < a_1 < \cdots < a_r$ such that there exists a section of L vanishing to that order at p. Alternatively, one can define $a_i = \max\{n : \ h^0(L(-n \cdot p)) \geqslant r + 1 - i$. The *ramification orders* are the nondecreasing sequence of integers $\alpha_0 \leqslant \alpha_1 \leqslant \cdots \leqslant \alpha_r$ defined by $\alpha_i = a_i - i$. We therefore make the following definition.

Definition 1.6. Let (C, p) be a smooth projective curve with a marked point. The variety of line bundles of rank at least r with ramification at least α at p is

$$W_d^{r,\alpha}(C,p) = \{L \in \mathrm{Pic}^d(C) : \ h^0(L(-(\alpha_i+i)p)) \geqslant r+1-i \text{ for } i = 0, 1, \cdots, r\}.$$

To the data g, r, d, α, we associate a partition $(\lambda_0, \cdots, \lambda_r)$ as follows (see Figure 1 for a visual illustration).

$$\lambda_i = (g - d + r) + \alpha_{r-i}$$

An alternative description that we find useful is that λ is the minimal partition whose Young diagram (regarded as a subset of $\mathbf{Z}^2_{>0}$) contains all of the points in

$$\{(g - d + r + \alpha_i, r + 1 - i) : \ i = 0, 1, \cdots, r\} \cap \mathbf{Z}^2_{>0}.$$

Observe that many such varieties $W_d^{r,\alpha}(C, p)$ are automatically isomorphic: if d is increased by 1, as are all of the integers α_i (that is, if a *base point* is added at p), the result is an isomorphic variety, now in $\mathrm{Pic}^{d+1}(C)$, obtained by applying the map $L \mapsto L(p)$ to the original. Therefore there is substantial redundancy in studying all possible data g, r, d, α. However, note that any two sets of data (g, r, d, α) that are isomorphic for the reason described give the same partition λ (the difference is simply where the gap is drawn in Figure 1). Indeed, we can identify both with a single subscheme of $\mathrm{Pic}^0(C)$, defined as follows.

Definition 1.7. The *Brill-Noether locus* corresponding to the partition λ and the marked curve (C, p) is

$$\begin{aligned} W^\lambda(C,p) = \{&L \in \mathrm{Pic}^0(C) : \ h^0(L(d' \cdot p)) \geqslant r' + 1 \\ &\text{whenever } (g - d' + r', r' + 1) \in \lambda\}. \end{aligned}$$

Remark 1.8. If $h^0(L(d \cdot p)) \geqslant r+1$, then it follows automatically that $h^0(L((d+1) \cdot p)) \geqslant r+1$ and $h^0(L((d-1) \cdot p)) \geqslant r$. In other words, once $(x, y) \in \lambda$, there are no additional restrictions imposed by $(x-1, y)$ or $(x, y-1)$. This explains why there is no loss of generality in requiring λ to be a partition, i.e. to be closed leftward and downward. Also, the Riemann-Roch theorem guarantees that if $r < 0$ or $g - d + r \leqslant 0$, then the condition $h^0(L(d \cdot p)) \geqslant r+1$ holds for *all* line bundles $L \in \mathrm{Pic}^0(C)$, so there is no loss of generality in only allowing λ to have points with both coordinates positive.

Comparing Definition 1.6 and Definition 1.7, it follows that the map $L \mapsto L(-d \cdot p)$ gives an isomorphism $W^{r,\alpha}_d(C, p) \to W^\lambda(C, p)$, where λ is the partition associated to g, r, d, α as described above. Note in particular that when λ is a rectangle (i.e. the minimal partition containing $(g-d+r, r+1)$), we have $W^\lambda(C, p) \cong W^r_d(C)$ (regardless of the marked point p). Therefore it suffices to study the geometry of $W^\lambda(C, p)$. We hope that the results of this paper will convince the reader that indexing Brill-Noether varieties (with one marked point) by the partition λ is the most natural way to study their geometry.

The scheme structure on $W^\lambda(C, p)$ is defined as an intersection of degeneracy loci of certain maps of vector bundles over $\mathrm{Pic}^0(C)$. We summarize this description in Section 4, where we also deduce from intersection theory that $\dim W^\lambda(C, p) \geqslant g - |\lambda|$, with equality if and only if the class of $W^\lambda(C, p)$ (either in the Chow group or in singular cohomology) is

$$[W^\lambda(C, p)] = \Theta^{|\lambda|} \Big/ \prod_{(x,y)\in\lambda} \mathrm{hook}(x, y). \tag{1}$$

In formula (1), $\mathrm{hook}(x, y)$ denotes the hook length of the box (x, y) in the Young diagram of λ, i.e. the total number of boxes in λ that are either of the form (x', y) with $x' \geqslant x$ or (x, y') with $y' \geqslant y$. In particular, since $\deg \Theta^g = g!$, this formula shows in the case $g = |\lambda|$ that if $W^\lambda(C, p)$ is 0-dimensional, then its cardinality is equal, by the hook length formula [FRT54], to the number of standard Young tableaux on λ. Observe in particular that equation (1) does not depend on g at all, but only on λ; we offer this as evidence that λ is the most natural combinatorial object with which to parameterize Brill-Noether varieties of marked curves.

Theorem 1.5 can therefore be restated to say that, for a general point (C, p) in $\mathcal{M}_{g,1}$ over an algebraically closed field, the dimension of $W^\lambda(C, p)$ is always equal to $g - |\lambda|$ (or $W^\lambda(C, p)$ is empty if $|\lambda| > g$), and formula (1) gives the class of this locus. This will be proved as Corollary 5.3.

1.2. **Chains of cycles.** If Γ is a metric graph, $w \in \Gamma$ is a point, and λ is a partition, then we define a generalized Brill-Noether locus $W^\lambda(\Gamma, w) \subseteq \mathrm{Pic}^0(\Gamma)$ in the same manner as for marked algebraic curves. Here $r(\cdot)$ denotes

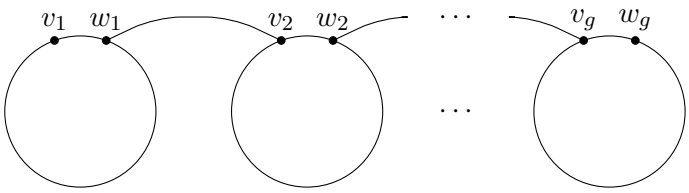

FIGURE 2. The chain of cycles Γ.

the Baker-Norine rank of a divisor (see [MZ08] or [GK08]).

$$W^\lambda(\Gamma, w) = \{[D] \in \mathrm{Pic}^0(\Gamma) : \ r(D + d' \cdot p) \geqslant r' \text{ whenever } (g - d' + r', r' + 1) \in \lambda\}$$

We will consider a metric graph Γ composed of a chain of cycles as shown in Figure 2. We require that v_i and w_i are *distinct* points for each i. The point w_g will serve as the marked point. We have drawn the chain with bridges from w_i to v_{i+1} to make some of our notation simpler, but all of our results are unaffected if the bridges have length 0.

We will show that all of the information about the Brill-Noether loci of (Γ, w_g) is encoded in a sequence $\underline{m} = (m_2, \cdots, m_g)$ of $g-1$ nonnegative integers, called the *torsion profile* of Γ. The integer m_i will be called the *ith torsion order* and is defined as follows.

Definition 1.9. Let ℓ_i denote the length of the ith cycle, and let $\ell(v_i w_i)$ denote the length of the clockwise edge from v_i to w_i. If $\ell(v_i w_i)$ is an irrational multiple of ℓ_i, then the ith torsion order m_i is 0. Otherwise, m_i is the minimum positive integer such that $m_i \cdot \ell(v_i w_i)$ is an integer multiple of ℓ_i.

An equivalent description (which explains the name) is that m_i identifies precisely which multiples of the divisor $w_i - v_i$ are linearly equivalent to 0. Note that $m_i \neq 1$, since we prohibit $v_i = w_i$.

We summarize this notation in the following situation, for easy reference.

Situation 1.10. Let Γ be a chain of cycles with points v_i, w_i as shown in Figure 2, and let m_i (for $i = 1, 2, \cdots, g$) denote the ith torsion order of this chain. The *torsion profile* of Γ is denoted by $\underline{m} = (m_2, \cdots, m_g)$.

Although we have defined m_i for all $i \in \{1, 2, \cdots, g\}$, note that we do not include the first torsion order m_1 in the torsion profile $\underline{m}$. This is because the value of m_1 is determined only by the position of v_1 on the first cycle, which in immaterial to the properties of the marked metric graph (Γ, w_g).

The Jacobian $\mathrm{Pic}^0(\Gamma)$ is naturally isomorphic to the product of all g cycles of Γ (see Lemma 3.3 for an explicit isomorphism that we will be the basis of our constructions). Our main result on the structure of $W^\lambda(\Gamma, w_g)$, Theorem 1.4, states that $W^\lambda(\Gamma, w_g)$ is always equal to a union of sub-tori, each given

by fixing a point on some of the cycles but not others. These tori are in bijection with $\underline{m}$-displacement tableaux, which depend on the torsion profile.

The proof of Theorem 1.4 is largely similar to the method in [CDPR12]: special divisors are constructed one cycle at a time, by placing a single chip on each cycle. The choice of where to place the chip of cycle i is determined in [CDPR12] by the data of a *lingering lattice path*, which can in turn be specified by a Young tableau on λ. We give a slightly different construction of divisors from Young tableaux. Unlike in [CDPR12], the divisors we produce are not necessary v_1-reduced divisors. However, the dependence of the chip positions on the tableau is more transparent. The idea underlying our construction is that we construct, inductively, divisors on a *sequence of chains* $\Gamma_1, \cdots, \Gamma_g = \Gamma$ of genera $1, 2, \cdots, g$, each obtained by taking only the first i cycles of Γ. Each divisor (when twisted by a multiple of w_i until it has degree 0) lies in a Brill-Noether locus $W^{\lambda_i}(\Gamma_i, w_i)$, where the partitions λ_i form an increasing sequence wherein the box (x, y) is added to the partition λ_i if and only if (x, y) has label i in the chosen tableau. Even in the case where we ultimately care only about the loci $W^r_d(\Gamma)$, the more general loci $W^\lambda(\Gamma, w_g)$ are essential, as they are used to record the intermediate steps in this construction. In this way, our method provides a geometric interpretation of the lingering lattice paths of [CDPR12]: a specific lingering lattice path records the precise way in which the ramification conditions grow as cycles are attached to Γ.

Remark 1.11. There are two contexts in algebraic geometry very similar to divisor theory on chains of cycles: limit linear series on chains of elliptic curves and limit linear series (as defined for curves of pseudo-compact type in [Oss16]) on stable curves formed as a chain of pairs of rational curves joined at two points. In both cases, we suspect that a similar combinatorial analysis should yield a description of the dimension of these spaces of limit linear series, and a procedure to enumerate their irreducible components. Osserman discusses the link between the latter context and that of tropical chains of cycles in [Oss16].

1.3. **Tropicalization.** The interplay between algebraic and tropical geometry will take place in the following situation.

Situation 1.12. In addition to the notation of Situation 1.10, let K be a complete valued field with infinite residue field, and let C be a smooth projective curve over K. Assume that C has totally split reduction and minimal skeleton isometric to Γ, and that w_g is rational over the value group of K. Let $p \in C$ be a K-point specializing to $w_g \in \Gamma$.

We show in Section 5 that Theorem 1.3 implies the following result on algebraic curves.

Theorem 1.13. *In Situation 1.12, if the torsion profile $\underline{m}$ of Γ satisfies $m_i = 0$ or $m_i > i$ for all i, then (C, p) is Brill-Noether general (as a marked algebraic curve).*

The generalized Brill-Noether Theorem 1.5 will follow, via Corollary 5.3.

1.4. **Organization of the paper.** We describe displacement tableaux in Section 2. The core of the paper is Section 3, in which we parameterize the Brill-Noether loci on marked chains of cycles and prove our main result, Theorem 1.4, which establishes a correspondence between the components of $W^\lambda(\Gamma, w_g)$ and displacement tableaux. We also deduce Theorems 1.1 and 1.3. Section 4 summarizes the basic facts about the scheme structure of $W^\lambda(C,p)$ in the algebraic case, and its analysis via intersection theory. Section 5 proves two semicontinuity properties needed to derive results in algebraic geometry from our tropical results, which we apply to prove Theorem 1.13 and complete our tropical proof of the generalized Brill-Noether Theorem 1.5. Finally, Section 6 states some further questions suggested by our work.

Acknowledgements. The key ideas for this paper came about through conversations with Melody Chan, David Jensen, and Sam Payne. The author is also grateful to Dhruv Ranganathan and Nicola Tarasca for several helpful conversations.

2. Displacement tableaux

The description in [CDPR12] of special divisors on generic chains gave a bijection between irreducible components of Brill-Noether loci and standard Young tableaux. Our description requires the following generalization of standard tableaux, which incorporates the possibility of special torsion profiles by allowing some symbols from the alphabet to occur more than once.

Definition 2.1. Let λ be a partition, and let $\underline{m} = (m_2, \cdots, m_g)$ be a $(g-1)$-tuple of nonnegative integers. An *$\underline{m}$-displacement tableux* on λ is a function $t: \lambda \to \{1, 2, \cdots, g\}$ satisfying the following properties.

(1) t is strictly increasing in any given row or column of λ.
(2) For any two distinct boxes $(x,y), (x',y')$ in λ such that $t(x,y) = t(x',y')$, the following congruence holds.

$$x - y \equiv x' - y' \pmod{m_{t(x,y)}}$$

We write $t \vdash_{\underline{m}} \lambda$ to denote that t is an $\underline{m}$-displacement tableau on λ.

Observe that if t satisfies property (1), then it is impossible for two different boxes in t to have the label 1. Therefore property (2) will never refer to the (unspecified) value of m_1.

Example 2.2. The tableau

6			
3	4	5	6
1	2	3	4

is a $(0,3,3,0,2)$-displacement tableau. The only repeated symbols are 3, 4, and 6, each of which occurs twice, in two places separated by a lattice distances 3, 3, and 4, respectively. Since $m_3 \mid 3$, $m_4 \mid 3$, and $m_6 \mid 4$, this is a $(0,3,3,0,2)$-displacement tableau.

Example 2.3. Let $\underline{0}$ denote the $(g-1)$-tuple $(0,0,\cdots,0)$. Then a $\underline{0}$-displacement tableau is the same thing as a standard Young tableau with alphabet $\{1,2,\cdots,g\}$ (note that the alphabet may have more elements than λ has boxes, so not all symbols from the alphabet need to appear).

The following notation will be convenient.

Notation 2.4. Let λ be a partition (identified, as usual, with a subset of $\mathbf{Z}_{>0} \times \mathbf{Z}_{>0}$). Let $\overline{\lambda}$ denote the union of λ with all points $(x,y) \in \mathbf{Z} \times \mathbf{Z}$ such that either $x \leqslant 0$ or $y \leqslant 0$.

Remark 2.5. In some sense, $\overline{\lambda}$ is the more natural set in our application, since the Riemann-Roch formula guarantees that if either coordinate of the point $(g-d+r, r+1)$ is nonpositive, then *all* degree d divisors have rank at least r.

A displacement tableau t may be interpreted as assembly instructions for the partition λ: boxes are added in g successive steps to produce a sequence of partitions. The permissible boxes to add at any given step are governed by the notion of displacement of a partition, which we now define (see Lemma 3.14 for the origin of displacement in our application).

Definition 2.6. Let λ be a partition, and let S be a set of integers. The *upward displacement* of λ by S is the partition

$$\mathrm{disp}^+(\lambda, S) = \lambda \cup L,$$

where L is the set of boxes $(x,y) \notin \lambda$ such that $x - y \in S$, $(x-1,y) \in \overline{\lambda}$, and $(x,y-1) \in \overline{\lambda}$. The elements of L are called the *loose boxes of λ with respect to S*.

The displacement process is easy to visualize: each element of $z \in S$ causes a slope-1 line through $(x,y) = (z,0)$ to be drawn through the Young diagram. Wherever these lines meet an inward corner, they turn that corner outwards, adding a box to the partition. Here, we regard the four corners of the box (x,y) as being the points $(x,y), (x-1,y), (x-1,y-1)$, and $(x,y-1)$ in the plane.

Example 2.7. Let $\lambda = (7,6,5,1)$, and let $S = 1 + 3\mathbf{Z}$. Then $\mathrm{disp}^+(\lambda, S) = (8,6,5,2)$, as illustrated in the figure below.

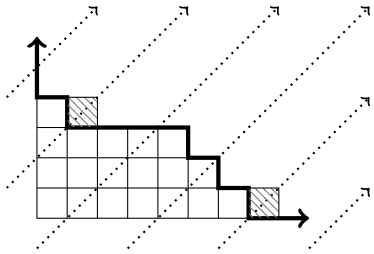

Lemma 2.8. *If λ, λ' are two partitions such that $\lambda \subseteq \lambda'$ and S is any set of integers, then* $\mathrm{disp}^+(\lambda, S) \subseteq \mathrm{disp}^+(\lambda', S)$.

Proof. Suppose that (x, y) is a box in $\mathrm{disp}^+(\lambda, S)$ that is not contained in λ'. Then certainly it is not contained in λ, so it was loose in λ. Therefore $(x-1, y)$ and $(x, y-1)$ both lie in $\overline{\lambda}$, and hence in $\overline{\lambda'}$. Thus (x, y) was also loose in λ', so it lies in $\mathrm{disp}^+(\lambda', S)$. Therefore every box in $\mathrm{disp}^+(\lambda, S)$ also lies in $\mathrm{disp}^+(\lambda', S)$. □

Displacement of partitions and displacement tableaux are related by the following combinatorial fact.

Corollary 2.9. *Let t be an $\underline{m}$-displacement tableau on a partition λ, and let $S_1, S_2, \cdots, S_g$ be sets of integers such that for all boxes $(x, y) \in \lambda$, $x - y \in S_{t(x,y)}$. Define a sequence of partitions $\lambda_0, \lambda_1, \cdots, \lambda_g$ inductively as follows.*

$$\lambda'_0 = \varnothing$$
$$\lambda'_{i+1} = \mathrm{disp}^+(\lambda'_i, S_{i+1})$$

Then $\lambda \subseteq \lambda'_g$.

Proof. Let $\lambda_i = t^{-1}(\{1, 2, \cdots, i\})$ be the sub-partition of λ given by boxes with labels less than or equal to i. By definition of displacement tableaux, all boxes with label $i+1$ are loose in λ_i with respect to the set S_i. Therefore $\lambda_{i+1} \subseteq \mathrm{disp}^+(\lambda_i, S_{i+1})$. It follows from Lemma 2.8 and induction on i that $\lambda_i \subseteq \lambda'_i$ for all i. □

3. Combinatorics of $W^\lambda(\Gamma, w_g)$

Let (Γ, w_g) be as in Situation 1.10. The language of displacement tableaux allows a succinct and explicit parameterization (Theorem 1.4) of all Brill-Noether loci $W^\lambda(\Gamma, w_g)$. We will use the following notation to conveniently denote points of the cycles of Γ.

Definition 3.1. In Situation 1.10, for any $\xi \in \mathbf{R}$ let $\langle \xi \rangle_i$ denote the point on the ith cycle that is located $\xi \cdot \ell(v_i w_i)$ units clockwise from w_i, where $\ell(v_i w_i)$ denotes the length of the clockwise edge from v_i to w_i.

Remark 3.2. In this notation,

(1) The points v_i and w_i are equal to $\langle -1 \rangle_i$ and $\langle 0 \rangle_i$, respectively;
(2) If n_1, n_2 are *integers*, then $\langle n_1 \rangle_i = \langle n_2 \rangle_i$ if and only if $n_1 \equiv n_2 \pmod{m_i}$, where m_i denotes the torsion order on the ith cycle.

The following lemma gives a convenient bijection between $\mathrm{Pic}^d(\Gamma)$ and the product of the g cycles of Γ.

Lemma 3.3. *Let D be any divisor of degree d on Γ. Then D is linearly equivalent to a unique divisor of the form*

$$\sum_{i=1}^{g} \langle \xi_i \rangle_i + (d-g) \cdot w_g.$$

Proof. Assume without loss of generality that $d = g$. So we must show that every degree g divisor is linearly equivalent to a unique sum of g points $\langle \xi_i \rangle_i$, with exactly one point on each cycle.

Existence of ξ_i: We may assume that the support of D is contained only on the cycles, since any point on a bridge is linearly equivalent to every other point on the the bridge. Next, by adding an integer linear combination of the divisors $w_i - v_{i+1}$ (each of which is linearly equivalent to 0), we may assume that the portion of D supported on the ith cycle has degree 1. By the Riemann-Roch formula, any degree 1 divisor on a genus 1 metric graph is linearly equivalent to an effective degree 1 divisor. Therefore we can replace the part of D supported on cycle i by a linearly equivalent divisor of the form $\langle \xi_i \rangle_i$ for some ξ_i.

Uniqueness of $\langle \xi_i \rangle_i$: Suppose that $D = \sum_{i=1}^{g} \langle \xi_i \rangle_i$ and $D' = \sum_{i=1}^{g} \langle \xi_i' \rangle_i$ are two linearly equivalent degree g divisors, each consisting of exactly one chip on each cycle. Then $D - D'$ is the principle divisor of a rational function f on Γ. Since the degree of the part of $D - D'$ to the left of v_i or the right of w_i is 0, the slope of f is 0 on all bridges. Therefore, by restricting f to the ith cycle, we obtain a rational function on the cycle whose principle divisor is equal to $\langle \xi_i \rangle_i - \langle \xi_i' \rangle_i$. But two points on a cycle are linearly equivalent as divisors if and only if they are the same point: otherwise the rational function whose principle divisor is their difference would have slope s on one interval, and slope $s+1$ on a complementary interval of the cycle, which is impossible since both these slopes would be nonnegative or both nonpositive, with at least one nonzero. □

Remark 3.4. The values ξ_i are easy to compute, given the divisor D. Define, for any point $p \in \Gamma$ and index i, the number $\widetilde{\xi}_i(p)$ to be -1 if p lies to the left of v_i, 0 if p lies to the right of w_i, and otherwise to be ζ such that $p = \langle \zeta \rangle_i$. The function $\widetilde{\xi}_i$ extends to all divisors by linearity. Then the desired values ξ_i are

$$\xi_i(D) = (i-1) + \widetilde{\xi}_i(D).$$

The value ξ_i is well-defined modulo the length of the the ith cycle. Note that $\xi_i(D)$ is not a linear function of the divisor D. Instead, it has the feature that $\xi_i(K_\Gamma - D) = -\xi_i(D)$, since $\xi_i(K_\Gamma) = 2(i-1)$.

Lemma 3.3 identifies $\mathrm{Pic}^d(\Gamma)$ with the product of the g cycles. Every displacement tableau $t \vdash_{\underline{m}} \lambda$ defines a sub-torus in $\mathrm{Pic}^0(\Gamma)$ as follows.

Definition 3.5. Use the notation of Situation 1.10. Let t be an $\underline{m}$-displacement tableau on a partition λ. Denote by $\mathbf{T}(t)$ the set of divisor classes of the form

$$\sum_{i=1}^{g} \langle \xi_i \rangle_i - g \cdot w_g,$$

where $\xi_1, \cdots, \xi_g$ are real numbers such that

$$\xi_{t(x,y)} \equiv x - y \mod m_{t(x,y)}$$

for all $(x, y) \in \lambda$.

Observe that while this definition makes use of the first torsion order m_1 of the chain in question, although this number is not included in the torsion profile $\underline{m}$. This apparent discrepancy is harmless since the value of m_1 is immaterial to the definition of an $\underline{m}$-displacement tableau.

Remark 3.6. If t^* is the dual tableau of t (obtained by switching the axes), then $\mathbf{T}(t^*)$ is the Serre dual of $\mathbf{T}(t)$, twisted by $-(2g-2)w_g$ (i.e. the set of classes $[K_\Gamma - D - (2g-2)w_g]$ such that $[D] \in \mathbf{T}(t)$).

The definition of $\underline{m}$-displacement tableaux ensures that $\mathbf{T}(t)$ is well-defined and non-empty, since no two boxes of t will impose conflicting conditions on ξ_i (due to the second part of Remark 3.2).

The main result of this paper, Theorem 1.4 from the introduction, asserts that

$$W^\lambda(\Gamma, w_g) = \bigcup_{t \vdash_{\underline{m}} \lambda} \mathbf{T}(t).$$

We prove Theorem 1.4 in Section 3.3. The key ingredient is Lemma 3.14, which illuminates the link between displacement of partitions and Brill-Noether loci.

3.1. **Applications and examples of Theorem 1.4.** For now, we will assume the result of Theorem 1.4, demonstrate how to use it in several examples, and deduce some consequences.

Corollary 3.7. *The dimension of the largest component of $W^\lambda(\Gamma, w_g)$ is equal to the maximum number of omitted symbols in an $\underline{m}$-displacement tableau on λ. This is equal to $g - |\lambda|$ if and only if it is impossible for an $\underline{m}$-displacement tableau on λ to have any repeated symbols.*

Proof. Each $\mathbf{T}(t)$ is homeomorphic to a torus, whose dimension is equal to the number of symbols $i \in \{1, 2, \cdots, g\}$ that do not occur in t. This dimension is equal to $g - |\lambda|$ if and only if each symbol in t occurs exactly once. □

Corollary 3.8. *The metric graph Γ is Brill-Noether general (as a metric graph without a marked point, i.e. in the original sense of* [CDPR12]*) if and only if every $\underline{m}$-displacement tableau on a* rectangular *partition has all*

entries distinct. The marked metric graph (Γ, w_g) *is Brill-Noether general (as a marked metric graph) if and only if every* $\underline{m}$*-displacement tableau on* any *partition has all entries distinct.*

Example 3.9. Suppose that $g = 4$ and consider the locus $W^1_3(\Gamma)$. This is in bijection with $W^\lambda(\Gamma, w_g)$, where $\lambda = (2, 2)$. From the theory of algebraic curves, we expect this to be finite set with two elements, unless Γ is special. Indeed, there are two standard Young tableaux on λ, which correspond to the following two divisor classes of degree 3.

3	4
1	2

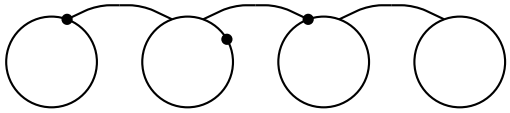

$\langle 0\rangle_1 + \langle 1\rangle_2 + \langle -1\rangle_3 + \langle 0\rangle_4 - w_g$

2	4
1	3

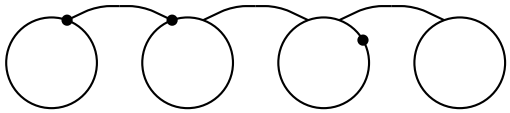

$\langle 0\rangle_1 + \langle -1\rangle_2 + \langle 1\rangle_3 + \langle 0\rangle_4 - w_g$

If $\underline{m} = \underline{0}$ (i.e. Γ is general), then these are the only points of $W^1_3(\Gamma)$. However, there are two situations in which additional points exist: if $m_2 = 2$ or if $m_3 = 2$. For example, if $m_2 = 2$, there are two additional one-dimensional tori in $W^1_3(\Gamma)$, coming from the two $(0, 2, 0, 0)$-displacement tableaux below. The second loop is drawn differently to show that $m_2 = 2$, and $*$s denote either variables or points which are free to vary (giving the degree of freedom of the torus).

2	4
1	2

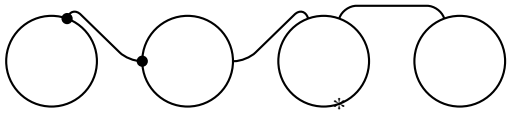

$\langle 0\rangle_1 + \langle 1\rangle_2 + \langle * \rangle_3 + \langle 0\rangle_4 - w_g$

2	3
1	2

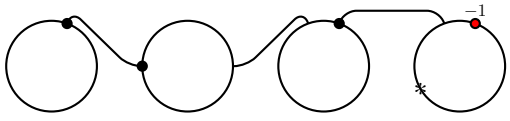

$\langle 0\rangle_1 + \langle 1\rangle_2 + \langle 0\rangle_3 + \langle * \rangle_4 - w_g$

Example 3.10. (*Hyperelliptic chains*) The locus $W^1_2(\Gamma)$ is nonempty if and only if there exists an $\underline{m}$-displacement tableau on $\lambda = (g - 1, g - 1)$. The

only possible displacement tableau on this partition is the following.

$$t = \begin{array}{|c|c|c|c|c|} \hline 2 & 3 & 4 & \cdots & g \\ \hline 1 & 2 & 3 & \cdots & g-1 \\ \hline \end{array}$$

This is an $\underline{m}$-displacement tableau if and only if $m_2 = m_3 = \cdots = m_{g-1} = 2$. Therefore a chain has a degree 2 divisor class of rank 1 if and only if all of these torsion orders are 2. Theorem 1.4 also shows that this divisor class is unique: it must be the divisor class of $\langle 0 \rangle_1 + \sum_{i=2}^{g} \langle i-3 \rangle_i - (g-2)w_g$.

We close this subsection by showing how Theorem 1.4 implies Theorems 1.1 and 1.3, which give necessary and sufficient criteria for a chain of cycles to be Brill-Noether general (in the sense of [CDPR12]), and for a marked chain of cycles (with marked point on the rightmost cycle) to be Brill-Noether general (as a marked metric graph).

Proof of Theorem 1.1. First, we show that the condition that $m_i = 0$ or $m_i > \min(i, g+1-i)$ is *necessary* for Brill-Noether generality. Suppose that i is an index such that $m_i \neq 0$, $m_i \leqslant g+1-i$, and $m_i \leqslant i$. Let $\lambda = (m_i, m_i)$ and define a function $t : \lambda \to \mathbf{Z}$ by

$$t(x,1) = x - m_i + i,$$
$$t(x,2) = x + i - 1.$$

The given inequalities guarantee that $1 \leqslant t(x,y) \leqslant g$, and t has exactly one repeated value: $t(m_i, 1) = t(1,2) = i$. Therefore t is an $\underline{m}$-displacement tableau on a rectangular partition with a repeated entry. It follows from Corollary 3.8 that Γ is not Brill-Noether general.

Now we show that the condition is *sufficient*. Suppose that Γ is *not* Brill-Noether general; we will show that there exists some i such that $0 < m_i \leqslant \min(i, g+1-i)$. By corollary 3.8, there exists a rectangular partition λ and a displacement tableau $t \vdash_{\underline{m}} \lambda$ with a repeated symbol. Select such a rectangular partition λ and tableau t so that $|\lambda|$ is as small as possible. Suppose that $(x,y), (x',y')$ are two distinct boxes containing the index i in t. Assume without loss of generality that $x < x'$ and $y > y'$. Then in fact (x,y) is the upper-left corner of λ, and (x',y') is the lower-right corner, since otherwise one can restrict t to the rectangle with corners $(x,y), (x,y'), (x',y'), (x',y)$ and regard this as an $\underline{m}$-displacement tableau on a smaller rectangle, contradicting the minimality of $|\lambda|$. It follows that $x = 1$, $y' = 1$, and the *only* repetition of symbols in t is $t(1,y) = t(x',1)$. Let this common label be i. Then the boxes $(1, y-1)$, $(1, y-2)$, $\cdots (1,1), (2,1), \cdots, (x'-1, 1)$ ($x'+y-3$ boxes total) all have distinct labels strictly less than i, so $x'+y-3 \leqslant i-1$. Since t is an $\underline{m}$-displacement tableau, $(y-1)-(1-x')$ is divisible by m_i; since $(y-1)-(1-x') = x'+y-2$ is positive, it follows that $m_i > 0$ and $m_i \leqslant x'+y-2 \leqslant i$. Similarly, all the boxes $(2,y), (3,y), \cdots, (x',y), (x',y-1), \cdots, (x',1)$ have distinct labels greater than i, and it follows that $x'+y-3 \leqslant g-i$, hence $m_i \leqslant g+1-i$. Therefore $0 < m_i \leqslant \min(i, g+1-i)$, as desired. $\square$

Proof of Theorem 1.3. First, suppose that there exists an index i such that $0 < m_i \leqslant i$. Let λ be the partition $(m_i, 1)$, and define a function t in exactly the same way as in the proof of Theorem 1.1 (but restricted to the smaller, non-rectangular partition). Then t is an $\underline{m}$-displacement tableau with the symbol i repeated, so (Γ, w_g) is not Brill-Noether general.

Conversely, suppose that (Γ, w_g) is not Brill-Noether general. We will show that there exists some $i \in \{2, 3 \cdots, g\}$ such that $m_i \leqslant i$. There exists some partition λ with a $\underline{m}$-displacement tableau t that repeats some symbol i. Suppose that $i = t(x, y) = t(x', y')$, and assume without loss of generality that $x < x'$ and $y > y'$. Then the boxes $(x, y-1), (x, y-2), \cdots, (x, y'), (x+1, y'), \cdots, (x'-1, y')$ constitute $(x-y)-(x'-y')-1$ boxes with *distinct* labels in t. Hence $i \geqslant (x-y)-(x'-y') > 0$. Since m_i divides $(x-y)-(x'-y')$, it follows that $m_i \leqslant i$. □

3.2. **Displacement of Weierstrass partitions.** We now begin the process of proving Theorem 1.4. We begin by linking the notion of displacement of partitions to divisors on metric graphs.

Definition 3.11. Let D be a degree 0 divisor on a marked metric graph (Γ, w). Define

$$\overline{\lambda}_{\Gamma,w}(D) = \{(g-d+r,\ r+1) :\ d, r \in \mathbf{Z},\ r(D+d\cdot w) \geqslant r\}.$$

If D is a divisor of *any* degree, define $\overline{\lambda}_{\Gamma,w}(D)$ to be $\overline{\lambda}_{\Gamma,w}(D - \deg D \cdot w)$. The *Weierstrass partition* of D at w is the set

$$\lambda_{\Gamma,w}(D) = \overline{\lambda}_{\Gamma,w}(D) \cap \mathbf{Z}^2_{>0}.$$

Remark 3.12. The word *Weierstrass* is chosen due to the following analogy with *Weierstrass semigroups* on algebraic curve: if (C, p) is a marked algebraic curve, D is a divisor, and $\lambda_{C,p}(D)$ is defined in the same manner as above, then $\lambda_{C,p}(0)$ (the Weierstrass partition of divisor 0) encodes the data of the Weierstrass semigroup of the point p.

This set $\lambda_{\Gamma,w}(D)$ is (the boxes in the Young diagram of) a partition since $r(D-w) \leqslant r(D)$ and $r(D+w) \leqslant r(D)+1$. The Riemann-Roch formula implies that $\overline{\lambda}_{\Gamma,w}$ always contains all points (x, y) with $x \leqslant 0$ or $y \leqslant 0$, so no information is lost in considering only the (finite) set $\lambda_{\Gamma,w}(D)$. The definition of $W^\lambda(\Gamma, w)$ implies that for any partition λ,

$$[D] \in W^\lambda(\Gamma, w) \text{ if and only if } \lambda \subseteq \lambda_{\Gamma,w}(D).$$

The purpose of this subsection and the next is to explain how to compute the Weierstrass partition of any divisor on a chain of cycles as in Situation 1.10. We begin in a more general situation. Let A be any metric graph, and $u \in A$ a point. Let C be a cycle, and v, w be two distinct points on C. Denote by Γ the metric graph obtained by attaching C to A by an edge connecting u to v (see Figure 3). We will explain how the Weierstrass partition of a divisor D at w on Γ is related to the Weierstrass partition of the restriction of D to the metric graph A at the marked point u. The

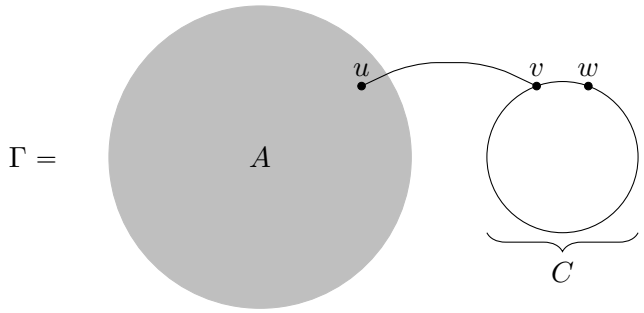

FIGURE 3. A marked point being displaced across a single cycle.

relationship is surprisingly simple, and it is the origin of our choice of the word displacement: when the marked graph (A, u) is "displaced" to the marked graph (Γ, w), the Weierstrass partition is displaced upward (in the sense of the disp^+ operation) in a manner than depends on the torsion order of the divisor $w - v$ (which in turn depends only on the edge lengths of C).

We will denote the points of C as follows, parallel to our notation when describing the chain of cycles. Here $\ell(vw)$ denotes the length of the clockwise edge from v to w.

$$\langle \xi \rangle = \text{the point on } C \text{ located } \xi \cdot \ell(vw) \text{ units clockwise from } w.$$

Let m denote the nonnegative generator of $\{n \in \mathbf{Z} : \langle n \rangle = \langle 0 \rangle\}$, i.e. the torsion order of the divisor $w - v$.

We summarize in the following lemma a few useful facts about metric graphs of this form.

Lemma 3.13. *Let Γ be the metric graph described above.*

(1) *Let D_A, D'_A be two divisors on A, and let D_C, D'_C be two divisors on C. Then $D_A + D_C$ is linearly equivalent to $D'_A + D'_C$ if and only if both*

$$D_A + \deg(D_C) \cdot u \sim D'_A + \deg(D'_C) \cdot u \text{ as divisors on } A, \text{ and}$$
$$D_C + \deg(D_A) \cdot v \sim D'_C + \deg(D'_A) \cdot v \text{ as divisors on } C.$$

(2) *Let $\ell(C)$ denote the length of the cycle C. Define a function $\sigma : \mathrm{Div}(C) \to \boldsymbol{R}/(\ell(C) \cdot \boldsymbol{Z})$ by $\langle \xi \rangle \mapsto (\xi + \ell(C) \cdot \boldsymbol{Z})$, extended by linearity. Two divisors on C are linearly equivalent if and only if they have the same degree and the same image under σ.*

(3) *Let D_A, D_C be divisors supported on A and C respectively. The linear series $|D_A + D_C|_\Gamma$ is nonempty if and only if there exists an integer n such that both $|D_A + n \cdot u|_A$ and $|D_C - n \cdot v|_C$ are nonempty. Here the subscript in the notation $|\cdot|_X$ indicates the metric graph on which the complete linear series is to be formed.*

Proof. The first part follows from the observation that a divisor D is principal on Γ if and only if the restriction of D to A differs from a principal

divisor by a multiple of u and the restriction of D to C differs from a principal divisor by a multiple of v. The second part is standard. The third part is a corollary of the first; we can take D'_A and D'_C to be the restrictions to A and C of any element of the complete linear series of $D_A + D_C$, after first moving all chips on the edge uv to one end or the other. □

The following lemma is the key to all of our analysis of chains of cycles.

Lemma 3.14. *Let D be any divisor of degree $g-1$ on A, where g is the genus of Γ, and p any point on C. Then*

$$\lambda_{\Gamma,w}(D+p) = \begin{cases} \mathrm{disp}^+(\lambda_{A,u}(D), z + m\boldsymbol{Z}) & \text{if } p = \langle z \rangle \text{ for some integer } z, \\ \lambda_{A,u}(D) & \text{otherwise.} \end{cases}$$

Proof. Throughout this proof, we will use the following vocabulary to make several statements more succinct: we call a divisor D on a marked metric graph (Γ, w) a *weak g^r_d for* (Γ, w) if $r(D+(d-\deg D)w) \geqslant r$. In other words, a weak g^r_d differs from a g^r_d by a multiple of the marked point[1].

Fix positive integers x, y, and define integers $r = y-1$ and $d = y-x+g-1$. Then $(x, y) \in \lambda_{\Gamma,w}(D+p)$ if and only if $D+p$ is a weak g^r_d for (Γ, w). We will establish the following two claims.

(1) If $p \neq \langle x-y \rangle$, then $D+p$ is a weak g^r_d for Γ if and only if D is a weak g^r_{d-1} for A.
(2) If $p = \langle x-y \rangle$, then $D+p$ is a weak g^r_d for Γ if and only if D is both a weak g^r_d and a weak g^{r-1}_{d-2} for A.

Translated into the language of Weierstrass partitions, these claims assert that $(x, y) \in \lambda_{\Gamma,w}(D+p)$ if and only if either:

(1') $p \neq \langle x-y \rangle$ and $(x, y) \in \lambda_{A,u}(D)$, or
(2') $p = \langle x-y \rangle$ and both $(x, y-1)$ and $(x-1, y)$ lie in $\overline{\lambda}_{A,u}(D)$.

We now show that these two claims imply the lemma. Consider first the case where $p = \langle z \rangle$ for some integer z. Then the set of all integers z' such that $\langle z' \rangle = \langle z \rangle$ is the arithmetic progression $z + m\mathbf{Z}$, hence $p = \langle x-y \rangle$ if any only if $x - y \equiv z \pmod m$. Therefore, claims (1) and (2) amount, in this case, to saying that $(x, y) \in \lambda_{\Gamma,w}(D+p)$ if and only if $(x, y) \in \mathrm{disp}^+(\lambda_{A,u}(D), z + m\mathbf{Z})$. Consider next the case where p is not equal to $\langle z \rangle$ for any integer z. Then $p \neq \langle x-y \rangle$ for all choices of (x, y), so claim (1) shows that $\lambda_{\Gamma,w}(D+p) = \lambda_{A,u}(D)$ in this case. Therefore the two claims will imply the lemma.

Proof of claims (1) and (2). Let ξ be a real number such that $p = \langle \xi \rangle$. Observe that the divisor $D+p+(d-g)w$ is linearly equivalent to $D+(d-g)v+\langle \xi+d-g \rangle$ by Lemma 3.13, which is equal to $D+(d-g)v+\langle \xi+y-x-1 \rangle$. Let D' denote this degree d divisor.

[1]In the context of algebraic curves, "a g^r_d" usually means a *linear series*, i.e. a choice of both a degree d divisor class and a $r+1$-dimensional space of sections of the associated line bundle. In our context, there is nothing analogous to a vector space of sections, so we will use the phrase to refer only to the divisor class.

By definition, $D+p$ is a weak g^r_d for (Γ, w) if and only if $|D'-E| \neq \varnothing$ for all effective divisors E of degree r. Replacing E by a linearly equivalent divisor if necessary, it suffices to consider only effective divisors E with at most one chip on C, and no chips on the interior of the edge from u to v.

Case 1: E has one chip on C and $r-1$ chips on A. Let the chip on C be placed at the point $\langle\delta\rangle$, and let $E_A = E - \langle\delta\rangle$ be the part supported on A. By Lemma 3.13, $|D'-E| \neq \varnothing$ if and only if there exists an integer n such that $|D-E_A+n\cdot u|_A \neq \varnothing$ and $|(d-g-n)\cdot v + \langle\xi+y-x-1\rangle - \langle\delta\rangle|_C \neq \varnothing$. By Lemma 3.13 and the fact that $v = \langle -1\rangle$, the second condition is equivalent to

$$|(d-g-n+1)\cdot v - \langle\delta-\xi-y+x\rangle| \neq \varnothing \text{ on } C.$$

This holds either if $d-g-n+1 \geqslant 2$ or if $d-g-n+1 = 1$ and $\langle\delta-\xi-y+x\rangle = v$. Therefore, as long as $\langle\delta\rangle$ is not equal to one specific point (namely, $\langle\xi+y-x-1\rangle$), the maximum value of n satisfying the second condition of Lemma 3.13(3) is $n = d-g-1$. Hence assuming that $\langle\delta\rangle$ is not equal to this point, $|D'-E| \neq \varnothing$ on Γ if and only if $|D+(d-g-1)u-E_A| \neq \varnothing$. This holds for *every* choice of degree-$(r-1)$ effective E_A on A if and only if D is a weak g^{r-1}_{d-2} for A.

We conclude that $|D+p-E| \neq \varnothing$ for *all* degree r effective E with exactly one chip on C if and only if D is a weak g^{r-1}_{d-2} on A.

Case 2: E is supported on A. Again, we analyze the maximum value of n such that when $n\cdot v$ is subtracted from the part of $D'-E$ supported on C, the result has nonempty linear series. The part of $D'-E$ supported on C consists of $(d-g)v + \langle\xi+y-x-1\rangle$, hence there are two cases: if $\langle\xi+y-x-1\rangle = v$, then the maximum n is $d-g+1$, and otherwise it is $d-g$. The hypothesis of the first case is equivalent to $p = \langle x-y\rangle$. Define

$$\epsilon = \begin{cases} 1 & \text{if } p = \langle x-y\rangle \\ 0 & \text{otherwise.} \end{cases}$$

Lemma 3.13 now says that $|D'-E| \neq \varnothing$ on Γ if and only if $|D+(d-g+\epsilon)u-E| \neq \varnothing$ on A. This holds for every effective degree r divisor on A if and only if D is a weak $g^r_{d-1+\epsilon}$ for A.

End of the proof: combining these two cases, we see that $D+p$ is a weak g^r_d for Γ if and only if D is both a weak g^{r-1}_{d-2} and a weak $g^r_{d-1+\epsilon}$ for A. If $p = \langle x-y\rangle$, i.e. $\epsilon = 1$, this is precisely claim (2). If $p \neq \langle x-y\rangle$, i.e. $\epsilon = 0$, then this is equivalent to saying only that D is a weak g^r_{d-1}, since a weak g^r_{d-1} is automatically also a weak g^{r-1}_{d-2}; this gives claim (1) and completes the proof of the lemma. □

3.3. **Proof of Theorem 1.4.** We now return Situation 1.10, where Γ is a chain of cycles with torsion profile $\underline{m}$. By induction on the genus (where the base case may be taken to be a genus 0 metric graph consisting of a single vertex), Lemma 3.14 immediately gives the a description, in terms

of displacement, of the Weierstrass partition of a divisor of the form $D = \sum_{i=1}^{g} \langle \xi_i \rangle_i$, as follows.

Corollary 3.15. *Let D be the degree g divisor $\sum_{i=1}^{g} \langle \xi \rangle_i$ on the marked graph (Γ, w_g) of Situation 1.10. For $i = 1, 2, \cdots, g$, let $S_i = \{z \in \boldsymbol{Z} : \langle \xi_i \rangle_i = \langle z \rangle_i\}$. This is either empty, a set with one element (in which case $m_i = 0$), or a congruence class modulo m_i. Define partitions $\lambda_0, \lambda_1, \cdots, \lambda_g$ as follows.*

$$\lambda_0 = \varnothing$$
$$\lambda_{i+1} = \mathrm{disp}^+(\lambda_i, S_{i+1})$$

Then the Weierstrass partition of D on (Γ, w_g) is λ_g.

Lemma 3.16. *If $t \vdash_{\underline{m}} \lambda$, then $\boldsymbol{T}(t) \subseteq W^{\lambda}(\Gamma, w_g)$.*

Proof. Suppose that $[D] \in \mathbf{T}(t)$. Define sets $S_1, S_2, \cdots, S_g$ and partitions $\lambda_0, \lambda_1, \cdots, \lambda_g$ as in corollary 3.15; by that corollary, λ_g is the Weierstrass partition of D for (Γ, w_g). It suffices to prove that $\lambda \subseteq \lambda_g$. By the definition of $\mathbf{T}(t)$, the sets S_i satisfy the condition that $x - y \in S_{t(x,y)}$ for all $(x, y) \in \lambda$. Therefore S_i also meet the hypothesis of Corollary 2.9, from which it follows that $\lambda \subseteq \lambda_g$, as desired. □

Lemma 3.17. *If $[D] \in W^{\lambda}(\Gamma, w_g)$, then there exists an $\underline{m}$-displacement tableau t such that $[D] \in \boldsymbol{T}(t)$.*

Proof. By Lemma 3.3, we may assume without loss of generality that $D + g \cdot w_g$ consists of exactly one chip on each cycle of Γ. Let the position of the chip on the ith cycle be $\langle \xi_i \rangle_i$.

For each $i \in \{1, 2, \cdots, g\}$, denote by Γ_i the first i cycles of the chain Γ (with the bridges between them), and denote by D_i the restriction of $D + g \cdot w_g$ to Γ_i. Let λ_i be the Weierstrass partition of D_i for the marked metric graph (Γ_i, w_i). Then by Lemma 3.14, these partitions are related by

$$\lambda_{i+1} = \mathrm{disp}^+(\lambda_i, S_{i+1}),$$

where $S_i = \{z \in \mathbf{Z} : \langle \xi_i \rangle_i = \langle z \rangle_i\}$. In particular, all of the boxes (x, y) in $\lambda_{i+1} \backslash \lambda_i$ have values of $x - y$ that are congruent modulo m_i.

Define a function $t : \lambda \to \{1, 2, \cdots, g\}$ by

$$t(x, y) = \min\{i : (x, y) \in \lambda_i\}.$$

The previous paragraph shows that t is in fact an $\underline{m}$-displacement tableau. Each $(x, y) \in \lambda$ is a loose box of $\lambda_{t(x,y)-1}$ with respect to the set $S_{t(x,y)}$, so in particular $x - y \in S_{t(x,y)}$, i.e. $\langle \xi_{t(x,y)} \rangle_{t(x,y)} = \langle x - y \rangle_{t(x,y)}$. Therefore $[D] \in \mathbf{T}(t)$. □

Theorem 1.4 follows directly from Lemmas 3.16 and 3.17.

4. Algebraic geometry of $W^\lambda(C,p)$

Throughout this section, let C denote a smooth projective curve over a field K, and $p \in C(K)$ a K-point. We will define the structure of $W^\lambda(C,p)$ as a subscheme of $\mathrm{Pic}^0(C)$, and compute its expected dimension and class via intersection theory.

We first review the scheme structure of $W^r_d(C)$, following [ACGH85, IV §3]. If an effective divisor E of degree at least $2g-d-1$ is fixed, then the vector spaces $H^0(L(E))$ (where L varies over all degree d line bundles) all have the same dimension $\deg E + d - g + 1$, and form the fibers of a vector bundle $\mathcal{M}$ over $\mathrm{Pic}^d(C)$. A second vector bundle $\mathcal{P}$ may be formed, whose fiber over $[L] \in \mathrm{Pic}^d(C)$ is $H^0(L(E)/L)$. There is an obvious map $\mathcal{M} \to \mathcal{P}$ of vector bundles. Observe that for any particular line bundle L of degree d, $h^0(L) \geqslant r+1$ if and only if $\dim \ker \left(H^0(L(E)) \to H^0(L(E)/L)\right) \geqslant (r+1)$. Therefore $W^r_d(C)$ is defined to be the subscheme of $\mathrm{Pic}^d(C)$ whose defining equations are, locally, the $N \times N$ minors of a trivialization of the bundle map $\mathcal{M} \to \mathcal{P}$, where $N = (\deg E + d - g - r + 1)$. The details of the construction of the vector bundles $\mathcal{M}, \mathcal{P}$, as well as the fact that this scheme does not depend on any of the choices made, can be found in [ACGH85].

Each scheme $W^r_d(C)$ may be regarded as a subscheme of $\mathrm{Pic}^0(C)$, via the isomorphism $L \mapsto L(-d \cdot p)$.

Definition 4.1. The scheme $W^\lambda(C,p)$ is the intersection of all the schemes $W^r_d(C)$ such that $(g-d+r, r+1) \in \lambda$, where these are regarded as subschemes of $\mathrm{Pic}^0(C)$.

Definition 4.1 is the most natural definition, but it is difficult to use since the various schemes being intersected are not transverse. We therefore give an alternative characterization, to which we can apply intersection theory.

First, since we wish to describe a locus in $\mathrm{Pic}^0(C)$ rather than $\mathrm{Pic}^d(C)$, we may replace L by $L(d \cdot p)$ in the discussion above.

Next, suppose that in the construction of the schemes $W^r_d(C)$ we choose the divisor E to be $(2g-1-d) \cdot p$. Then the same vector bundle $\mathcal{M}$ can be used in the construction of every locus $W^r_d(C)$; its fiber over $[L]$ will be $H^0(L((2g-1)p))$. The bundle $\mathcal{P}$ used in the definition of $W^r_d(C)$ has fibers naturally identified with $H^0(L((2g-1)\cdot p)/L(d \cdot p))$ (note that we now write $L(d \cdot p)$ since L is a degree 0 line bundle, and $L(d \cdot p)$ is the line bundle which is required to have $r+1$ sections). Write this bundle as $\mathcal{P}_d$. These vector bundles, for various choices of d, fit into a flag of quotients.

$$\mathcal{P} = \mathcal{P}_{-1} \twoheadrightarrow \mathcal{P}_0 \twoheadrightarrow \mathcal{P}_1 \twoheadrightarrow \cdots \twoheadrightarrow \mathcal{P}_{2g-1} = 0$$

In this notation, we can alternatively characterize $W^\lambda(C,p)$ as the following degeneracy locus.

$$W^\lambda(C,p) = \big\{[L] \in \mathrm{Pic}^0(C) : \ \mathrm{rank}(\mathcal{M}_{[L]} \to (\mathcal{P}_d)_{[L]}) \leqslant g-r-1 \\ \text{whenever } (r+1, g-d+r) \in \lambda\}$$

Degeneracy loci of this form are analyzed in [Ful92]. The expected codimension of this degeneracy locus is equal to $|\lambda|$, and the class of the locus, in case it has the correct dimension, can be computed from the Chern classes of $\mathcal{M}$ and $\mathcal{P}_d$. By [ACGH85, VII §4] (where the same bundles are used to analyze $W^r_d(C)$), $\mathcal{P}_i$ has trivial Chern classes, and $\mathcal{M}$ has Chern classes given by

$$c_i(-\mathcal{M}) = \Theta^i/i!.$$

By [Ful92, Theorem 10.1], this degeneracy locus supports a (Chow or singular cohomology) class given by the following determinant, with this class being equal to the class of the degeneracy locus in case the dimensions match. Here λ_1^* denotes the largest element of the dual partition, i.e. the number of nonzero elements of λ.

$$\begin{aligned}\Omega^\lambda(C) &= \det\left(c_{\lambda_i - i + j}(-\mathcal{M})\right)_{1 \leqslant i,j \leqslant \lambda_1^*} \\ &= \det\left(\Theta^{\lambda_i - i + j}/(\lambda_i - i + j)!\right)_{1 \leqslant i,j \leqslant \lambda_1^*} \\ &= \Theta^{|\lambda|} \cdot \det\left(1/(\lambda_i - i + j)!\right)_{1 \leqslant i,j \leqslant \lambda_1^*}\end{aligned}$$

In the last line, the number $1/n!$ should be interpreted to be 0 when $n < 0$. By Aitken's determinantal formula [Ait43], this determinant is equal to the number of standard Young tableaux on λ divided by $|\lambda|!$; by the hook length formula [FRT54] we can express this as the product of the reciprocals of the hook lengths of λ.

We therefore deduce the following fact.

Proposition 4.2. *If $g \geqslant |\lambda|$, then the Brill-Noether locus $W^\lambda(C,p)$ is nonempty of dimension at least $g - |\lambda|$. If the dimension of $W^\lambda(C,p)$ is exactly $g - |\lambda|$, then its (Chow or singular cohomology) class is equal to*

$$\Theta^{|\lambda|} \Big/ \prod_{(x,y) \in \lambda} \mathrm{hook}(x,y) \ .$$

5. Tropicalization

This section is concerned with the marked curve (C,p) from Situation 1.12. We will prove Theorem 1.13, which is a corollary of a pair of "semicontinuity" results.

The retraction from the Berkovich analytic space C^{an} to its skeleton extends by linearity to divisors and induces a map

$$\mathrm{Trop} : \ \mathrm{Pic}(C) \to \mathrm{Pic}(\Gamma).$$

When K is the field of fractions of a discrete valuation ring, this map is identical to the tropicalization map of [Bak08]. This map has two important properties, which can be interpreted as two forms of semicontinuity for Trop.

Proposition 5.1.

(1) *The image of* $W^\lambda(C,p)$ *under* Trop *lies in the locus* $W^\lambda(\Gamma, w_g)$.
(2) $\dim W^\lambda(\Gamma, w_g) \geqslant \dim W^\lambda(C,p)$.

Proof. The definition of Trop implies that for any divisor D on C and any integer n, $\mathrm{Trop}(D + n \cdot p) = \mathrm{Trop}(D) + nw_g$. Therefore it suffices to show that $\mathrm{Trop}(W^r_d(C)) \subseteq W^r_d(\Gamma)$ for all integers r, d. This is precisely the specialization inequality of [Bak08]; although that paper works in the discrete valuation case, the proof applies without modification to the general case. This establishes part (1).

The tropical Jacobian $\mathrm{Pic}^0(\Gamma)$ is isomorphic to the skeleton $\Sigma(\mathrm{Pic}^0(C)^{\mathrm{an}})$ [BR15], and this isomorphism is compatible with the tropicalization map. Therefore Gubler's results on tropicalization of subvarieties of abelian varieties apply: by [Gub07, Theorem 6.9], a variety of pure dimension d in $\mathrm{Pic}^0(C)$ has image, under the retraction to the skeleton $\Sigma(\mathrm{Pic}^0(C))$, of pure dimension d. Combining this with part (1), it follows that the image $\mathrm{Trop}(W^\lambda(C,p)) \subseteq \Gamma$ contains a tropical variety of dimension equal to the dimension of $W^\lambda(C,p)$, and part (2) follows. □

Corollary 5.2 (Theorem 1.13)**.** *In Situation 1.12, if for each* i *either* $m_i = 0$ *or* $m_i > i$*, then for every partition* λ*,* $W^\lambda(C,p)$ *has the expected dimension.*

Proof. Theorem 1.3 and Proposition 5.1. □

Corollary 5.3. *Working over any algebraically closed field, if* (C,p) *is a general marked curve, then for every partition* λ*,* $W^\lambda(C,p)$ *has the expected dimension, and Chow class given by Formula* (1)*.*

Proof. Theorem 1.13 and deformation theory (cf. [Bak08, Appendix]) show the existence of Brill-Noether general marked curves over any complete valued field with infinite residue field. The existence of Brill-Noether general marked curves over an arbitrary algebraically closed field follows from the fact that the coarse moduli space $M_{g,1}$ of marked smooth curves is a scheme of finite type over Spec$\mathbf{Z}$, and the locus of Brill-Noether general points is Zariski open and surjects onto Spec$\mathbf{Z}$.

Proposition 4.2 shows that the Chow class is given by formula (1). □

The generalized Brill-Noether theorem (Theorem 1.5) follows from Corollary 5.3.

6. Further questions for arbitrary metric graphs

The detailed description, for marked chains of cycles (Γ, w_g), of the Brill-Noether loci $W^\lambda(\Gamma, w_g)$, makes it possible to see an intriguing analogy with the theory of algebraic curves; we believe that extending this analogy to

arbitrary metric graphs might yield further insight. In this section we formulate this analogy and the questions arising from it.

For a marked algebraic curve (C, p), classical intersection theory, as discussed in Section 4, guarantees not only that $W^\lambda(C, p)$ is nonempty (when $g - |\lambda| \geqslant 0$), but that all of its irreducible components are at least $(g - |\lambda|)$-dimensional and that it supports a particular intersection class, namely $\Theta_C^{|\lambda|} / \prod_{(x,y)\in\lambda} \mathrm{hook}(x, y)$. Furthermore, this is precisely the class of $W^\lambda(C, p)$ when its dimension is exactly $g - |\lambda|$.

Even though there is currently no equivalent in tropical intersection theory of the machinery available to study degeneracy loci in algebraic geometry (as in [Ful92]), the same facts outlined in the previous paragraph hold in the context of chains of cycles, when interpreted in a suitable way. Indeed, observe that if (Γ, w_g) is a marked chain of cycles as in Situation 1.10, and λ is a partition, then regardless of the torsion profile there is always a locus given by those tableaux that have no repeating symbols, namely

$$\bigcup_{t \vdash \underline{0}\lambda} \mathbf{T}(t) \subseteq W^\lambda(\Gamma, w_g).$$

The parameterization provided by Defintion 3.5 shows that this locus varies continuously as the edge lengths of Γ vary. In this section we will refer to this subset as the *stable locus*. It consists of a union of $(g - |\lambda|)$-dimensional tori. Furthermore, observe that any point of $W^\lambda(\Gamma, w_g)$ that is not in the stable locus necessarily sits in a torus $\mathbf{T}(t)$ coming from a tableau t with a repeated symbol, i.e. a torus of dimension strictly greater than $g - |\lambda|$. Thus the locus $W^\lambda(\Gamma, w_g)$ is equal to the stable locus if and only if it is exactly $(g - |\lambda|)$-dimensional.

To complete the analogy with algebraic curves, we point out that, in a suitable sense, the stable locus has intersection class $\Theta_\Gamma^{|\lambda|} / \prod_{(x,y)\in\lambda} \mathrm{hook}(x, y)$. We must define what we mean by these terms.

For any metric graph Γ of genus g, there is a tropical subvariety $\Theta_\Gamma \subset \mathrm{Pic}^{g-1}$, called the tropical theta divisor, which can be identified with $W^0_{g-1}(\Gamma)$; see [MZ08]. As usual, we suppose that we have chosen a marked point $w \in \Gamma$, and regard Θ_Γ as a subset of $\mathrm{Pic}^0(\Gamma)$. When Γ is a chain of cycles, Theorem 1.4 shows that Θ_Γ is a union of $(g-1)$-dimensional tori (the Young diagram associated to $r = 0,\ d = g - 1$ is a single box), each obtained by fixing one coordinate ξ_i to be equal to 0 (in the notation of Lemma 3.3). If d is a positive integer, and we choose d general translates of Θ_Γ, we can easily describe the intersection of these translates: it is a union of $g(g-1)\cdots(g-d+1)$ sub-tori, each giving by fixing d coordinates ξ_i. More precisely, for each choice of d coordinates, the intersection includes $d!$ tori given by fixing these d coordinates.

Now, observe that the stable locus of $W^\lambda(\Gamma, w_g)$ has a similar description: for each standard Young tableaux on λ and choice of $|\lambda|$ symbols from the set $\{1, 2, \cdots, g\}$, there is one torus in the stable locus, given by fixing the

$|\lambda|$ coordinates chosen. By the hook-length formula [FRT54], the number of tori in the stable locus fixing a given set of $|\lambda|$ coordinates is

$$|\lambda|!/\prod_{(x,y)\in\lambda}\text{hook}(x,y).$$

Provided that any subtorus of $\text{Pic}^0(\Gamma)$ is regarded as "equivalent" to any translate of itself, it follows that the stable locus of $W^\lambda(\Gamma, w_g)$ may be regarded as being equivalent to $1/\prod_{(x,y)\in\lambda}\text{hook}(x,y)$ times the locus given by intersecting $|\lambda|$ general translates of Θ_Γ. This precisely mirrors the formula for the expected class of $W^\lambda(C,p)$ in the algebraic case.

To formulate a more general statement for arbitrary marked metric graphs (Γ, w), we require a suitable equivalence relation on tropical subvarieties of $\text{Pic}^0(\Gamma)$. One candidate is the notion of rational equivalence described in [AHR16]. Whatever choice is made, it should be compatible with tropical intersection theory, e.g. as developed in [AR10].

To extend the analogy outlined above to arbitrary marked metric graphs, several questions present themselves.

Question 6.1. If (Γ, w) is any marked metric graph, is the local dimension of $W^\lambda(\Gamma, w)$ at least $g - |\lambda|$ at every point?

Note that the semicontinuity results of Section 5 imply that the local dimension of $W^\lambda(\Gamma, w)$ is at least $g - |\lambda|$ at every point that lifts to a point of $W^\lambda(C,p)$ on an algebraic marked curve. This does not rule out the existence of extra components of smaller dimension.

Question 6.2. If (Γ, w) is any marked metric graph, does the locus $W^\lambda(\Gamma, w)$ always support a tropical subvariety equivalent (in a suitable sense) to $\Theta_\Gamma^{|\lambda|}/\prod_{(x,y)\in\lambda}\text{hook}(x,y)$? Is there a "canonical" such locus that varies continuously as the edge lengths vary (analogous to the "stable" locus defined above for chains of cycles)?

Question 6.3. If Γ is a metric graph, w is a marked point, and λ is a partition such that $W^\lambda(\Gamma, w)$ has pure dimension $g-|\lambda|$, must $W^\lambda(\Gamma, w)$ be equivalent (in a suitable sense) to $\Theta_\Gamma^{|\lambda|}/\prod_{(x,y)\in\lambda}\text{hook}(x,y)$?

References

[ACGH85] E. Arbarello, M. Cornalba, P. A. Griffiths, and J. Harris. *Geometry of algebraic curves. Vol. I*, volume 267 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, New York, 1985.

[AHR16] Lars Allermann, Simon Hampe, and Johannes Rau. On rational equivalence in tropical geometry. *Canad. J. Math.*, 68(2):241–257, 2016.

[Ait43] A. C. Aitken. The monomial expansion of determinantal symmetric functions. *Proc. Roy. Soc. Edinburgh. Sect. A.*, 61:300–310, 1943.

[AR10] Lars Allermann and Johannes Rau. First steps in tropical intersection theory. *Math. Z.*, 264(3):633–670, 2010.

[Bak08] Matthew and Baker. Specialization of linear systems from curves to graphs. *Algebra Number Theory*, 2(6):613–653, 2008. With an appendix by Brian Conrad.

[BR15] Matthew Baker and Joseph Rabinoff. The skeleton of the Jacobian, the Jacobian of the skeleton, and lifting meromorphic functions from tropical to algebraic curves. *Int. Math. Res. Not. IMRN*, (16):7436–7472, 2015.

[CDPR12] Filip Cools, Jan Draisma, Sam Payne, and Elina Robeva. A tropical proof of the Brill-Noether theorem. *Adv. Math.*, 230(2):759–776, 2012.

[EH86] David Eisenbud and Joe Harris. Limit linear series: basic theory. *Invent. Math.*, 85(2):337–371, 1986.

[FRT54] J. S. Frame, G. de B. Robinson, and R. M. Thrall. The hook graphs of the symmetric groups. *Canadian J. Math.*, 6:316–324, 1954.

[Ful92] William Fulton. Flags, Schubert polynomials, degeneracy loci, and determinantal formulas. *Duke Math. J.*, 65(3):381–420, 1992.

[GK08] Andreas Gathmann and Michael Kerber. A Riemann-Roch theorem in tropical geometry. *Math. Z.*, 259(1):217–230, 2008.

[Gub07] Walter Gubler. Tropical varieties for non-Archimedean analytic spaces. *Invent. Math.*, 169(2):321–376, 2007.

[MZ08] Grigory Mikhalkin and Ilia Zharkov. Tropical curves, their Jacobians and theta functions. In *Curves and abelian varieties*, volume 465 of *Contemp. Math.*, pages 203–230. Amer. Math. Soc., Providence, RI, 2008.

[Oss16] Brian Osserman. Dimension counts for limit linear series on curves not of compact type. *Math. Z.*, 284(1-2):69–93, 2016.

[Pfl16] Nathan Pflueger. Brill-Noether varieties of k-gonal curves. *arXiv preprint arXiv:1603.08856*, 2016.

Department of Mathematics, Brown University, Box 1917, Providence, RI 02912

E-mail address: pflueger@math.brown.edu